arXiv:1609.07038v1 [cs.MA] 22 Sep 2016

Simultaneous Intermittent Communication Control and Path Optimization in Networks of Mobile Robots

Yiannis Kantaros and Michael M. Zavlanos

***Abstract*— In this paper, we propose an intermittent communication framework for mobile robot networks. Specifically, we consider robots that move along the edges of a connected mobility graph and communicate only when they meet at the nodes of that graph giving rise to a dynamic communication network. Our proposed distributed controllers ensure intermittent connectivity of the network and path optimization, simultaneously. We show that the intermittent connectivity requirement can be encapsulated by a global Linear Temporal Logic (LTL) formula. Then we approximately decompose it into local LTL expressions which are then assigned to the robots. To avoid conflicting robot behaviors that can occur due to this approximate decomposition, we develop a distributed conflict resolution scheme that generates non-conflicting discrete motion plans for every robot, based on the assigned local LTL expressions, whose composition satisfies the global LTL formula. By appropriately introducing delays in the execution of the generated motion plans we also show that the proposed controllers can be executed asynchronously.**

I. Introduction

Communication among robots has been typically modeled using proximity graphs and the communication problem is often treated as preservation of graph connectivity [1]–[4]. Common in the above works is that point-to-point or end-to-end network connectivity is required to be preserved for all time. However, this requirement is often very conservative, since limited resources may hinder robots from accomplishing their assigned goals. Motivated by this fact, in this paper we propose a *distributed intermittent communication protocol* for mobile networks. In particular, we consider that robots move along the edges of a connected mobility graph and communicate only when they meet at the nodes of this graph giving rise to a dynamic communication network. We design distributed controllers that ensure intermittent communication of the network while minimizing at the same time the distance traveled between meeting points. We show that intermittent communication can be captured by a global Linear Temporal Logic (LTL) formula that forces robots to meet infinitely often at the rendezvous points. Given such an LTL expression, existing model checking techniques [5], [6] can be employed in order to implement correct by construction controllers for all robots.

LTL-based control synthesis and task specification for mobile robots build upon either a bottom-up approach when independent LTL expressions are assigned to robots [7]–[9] or top-down approaches when a global LTL describing a collaborative task is assigned to a team of robots [10], [11], as in our work. Bottom-up approaches generate a discrete high-level motion plan for all robots based on a synchronous product automaton among all agents and, therefore, they are resource demanding and scale poorly with the number of robots. To mitigate these issues, we propose a novel technique that approximately decomposes the global LTL formula into local ones and assigns them to robots. Since the approximate decomposition of the global LTL formula can result in conflicting robot behaviors we develop a distributed conflict resolution scheme that generates discrete motion plans for every robot based on the assigned local LTL expressions. By appropriately introducing delays in the execution of the generated motion plans we show the proposed controllers can also be executed in an asynchronous fashion. In contrast, most relevant literature assumes that robot control is performed in a synchronous way [10], [11]. Asynchronous robot mobility is considered in [12] by introducing 'traveling states' in the discretized abstraction of the environment decreasing in this way the scalability of the proposed algorithm.

The most relevant works to the one proposed here are presented in [13]–[16]. A centralized intermittent communication control scheme is presented in [13] that ensures communication among robots infinitely often but it does not scale well with the number of robots. In [14] a distributed intermittent communication control scheme is proposed that requires synchronization among robots, unlike the approach developed here. Moreover, [14] considers *a priori* determined communication points while in this paper, rendezvous points are selected optimally, to minimize the distance traveled by the robots. [15] proposes a distributed synchronization scheme that allows robots that move along the edges of a bipartite mobility graph to meet periodically at the vertices of this graph. Instead, here we make no assumptions on the graph structure on which robots reside or on the communication pattern to be achieved. On the other hand, [16] proposes a receding horizon framework for periodic connectivity that ensures recovery of end-to-end connectivity within a given time horizon. As the number of robots or the size of the time horizon grows, this approach can become computationally expensive. To the contrary, our proposed method scales well to large numbers of robots and can handle situations where the whole network can not be connected at once, by ensuring connectivity over time, infinitely often.

This work is supported in part by the NSF awards CNS #1261828 and CNS #1302284.

Yiannis Kantaros and Michael M. Zavlanos are with the Department of Mechanical Engineering and Materials Science, Duke University, Durham, NC 27708, USA. {yiannis.kantaros,michael.zavlanos}@duke.edu

II. Problem Formulation

Consider a team of N robots that move in a workspace $\mathcal{W} \subset \mathbb{R}^n$ according to $\dot{\mathbf{x}}_i(t) = \mathbf{u}_i(t)$, where $\mathbf{x}_i(t) \in \mathbb{R}^n$ is the position of robot i, $i \in \{1, 2, \ldots, N\}$, at time t and $\mathbf{u}_i(t) \in \mathbb{R}^n$ is a control input. Also, consider L locations in $\mathcal{W}$ denoted by ℓ_j, $j \in \{1, 2 \ldots, L\}$ located at positions $\boldsymbol{\mu}_j \in \mathcal{W}$ and paths $\gamma_{ij} : [0, 1] \to \mathbb{R}^n$ that connect two locations ℓ_i and ℓ_j such that $\gamma_{ij}(0) = \boldsymbol{\mu}_i$ and $\gamma_{ij}(1) = \boldsymbol{\mu}_j$. The union of locations $\boldsymbol{\mu}_j$ and paths γ_{ij} gives rise to an undirected graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, where the set of nodes $\mathcal{V} = \{1, 2, \ldots, L\}$ is indexed by the set of locations ℓ_j and the set of edges $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is determined by the paths γ_{ij} such that an edge $(i, j) \in \mathcal{E}$ exists if and only if a path γ_{ij} exists. Hereafter, we call $\mathcal{G}$ a mobility graph and we assume that the robots move along its edges to possibly accomplish a high-level task; see Figure 1(a). For example, robots may travel along their assigned paths to monitor different parts of a region and then coordinate to transmit their measurements to a user. The presence of a network allows the robots to transmit data in a multi-hop fashion, so that they do not have to leave their assigned region. Applications of this framework involve distributed coverage, state estimation, or surveillance. In what follows, we assume that the mobility graph $\mathcal{G}$ is connected.

A. Intermittent Communication

We assume that the robotic team is divided into M subgroups. The indices i of robots that belong to the m-th subgroup are collected in a set denoted by $\mathcal{T}_m$, $m \in \{1, 2, \ldots, M\}$ while every robot can belong to more than one subgroups. Robots in a subgroup $\mathcal{T}_m$ can communicate only when all of them are present simultaneously at a common location ℓ_j. The locations ℓ_j where communication can take place for the robotic team $\mathcal{T}_m$ are collected in a set $\mathcal{C}_m$. This way, a dynamic robot communication graph $\mathcal{G}_c = \{\mathcal{V}_c, \mathcal{E}_c\}$ is constructed where the set of nodes $\mathcal{V}_c$ is indexed by robots, i.e., $\mathcal{V}_c = \{1, 2, \ldots, N\}$ and $\mathcal{E}_c \subseteq \mathcal{V}_c \times \mathcal{V}_c$ is the set of communication links that emerge when, e.g., all robots in a team $\mathcal{T}_m$ meet at a common rendezvous point $\ell_j \in \mathcal{C}_m$ simultaneously. Given the robot teams $\mathcal{T}_m$, a graph $\mathcal{G}_\mathcal{T} = \{\mathcal{V}_\mathcal{T}, \mathcal{E}_\mathcal{T}\}$ is constructed whose set of nodes $\mathcal{V}_\mathcal{T} = \{1, 2, \ldots, M\}$ is indexed by the teams $\mathcal{T}_m$ and set of edges $\mathcal{E}_\mathcal{T}$ consists of links between nodes m and n if $\mathcal{T}_m \cap \mathcal{T}_n \neq \varnothing$; see Figure 1(b). We assume that the graph $\mathcal{G}_\mathcal{T}$ is connected in order to ensure dissemination of information in the network. Furthermore, we can define the set of neighbors of node $m \in \mathcal{V}_\mathcal{T}$ by $\mathcal{N}_{\mathcal{T}_m} = \{n \in \mathcal{V}_\mathcal{T} | (n, m) \in \mathcal{E}_\mathcal{T}\}$. Also, for robot i we define the set that collects the indices of all other robots that belong to common teams with robot i as $\mathcal{N}_i = \{j | j \in \mathcal{T}_m, \forall m \in \mathcal{S}_i\} \setminus \{i\}$, and the set that collects the indices of teams that robot i belongs to as $\mathcal{S}_i = \{m | i \in \mathcal{T}_m,\ m \in \{1, 2, \ldots, M\}\}$.

At a rendezvous point $\ell_j \in \mathcal{C}_m$ communication takes place when all robots in $\mathcal{T}_m$ are present there, simultaneously. Then the communication graph $\mathcal{G}_c$ is defined to be *connected over time* if all robots in $\mathcal{T}_m$ meet infinitely often at a region

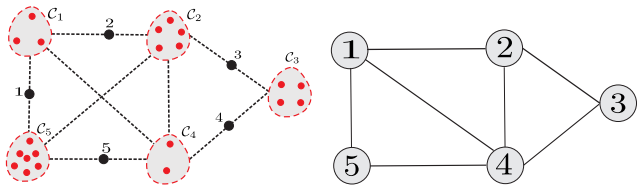

(a) Mobility Graph $\mathcal{G}$ (b) Graph $\mathcal{G}_\mathcal{T}$

Fig. 1. A graphical illustration of the problem formulation. Figure 1(a) depicts the mobility graph $\mathcal{G}$ for a network of $N = 5$ robots (black dots) divided into $M = 5$ teams. The robot teams are selected to be: $\mathcal{T}_1 = \{1, 2\}$, $\mathcal{T}_2 = \{2, 3\}$, $\mathcal{T}_3 = \{3, 4\}$, $\mathcal{T}_4 = \{2, 4, 5\}$, and $\mathcal{T}_5 = \{1, 5\}$. Red dots represent communication points and the black dashed lines between two sets $\mathcal{C}_m$ and $\mathcal{C}_n$ imply that there is at least one path γ_{ij} that connects communication points $\ell_i \in \mathcal{C}_m$ and $\ell_j \in \mathcal{C}_n$. Figure 1(b) depicts the associated graph $\mathcal{G}_\mathcal{T}$.

$\ell_j \in \mathcal{C}_m$, for all $m \in \mathcal{V}_\mathcal{T}$. Such a requirement can be captured by the following global LTL expression:

$$\phi = \wedge_{\forall m \in \{1,2,\ldots,M\}} \left(\square \Diamond \left(\vee_{\ell_j \in \mathcal{C}_m} (\wedge_{\forall i \in \mathcal{T}_m} \pi_i^{\ell_j}) \right) \right), \quad (1)$$

where $\pi_i^{\ell_j}$ is an atomic proposition that is true when robot i is sufficiently close to $\boldsymbol{\mu}_j$. In (1), $\wedge$ and $\vee$ are the conjunction and disjunction operator, respectively, while $\square$ and $\Diamond$ stand for the temporal operators 'always' and 'eventually', respectively. For more details on LTL, we refer the reader to [5], [6]. In words, the LTL expression in (1) requires all robots in a team $\mathcal{T}_m$ to meet infinitely often at at least one communication point $\ell_j \in \mathcal{C}_m$, for all $m \in \mathcal{V}_\mathcal{T}$.

B. Discretized Abstraction of the Workspace

We can model the environment in which robot i resides by a weighted transition system (wTS) denoted by wTS_i that is defined as follows

Definition 2.1 (weighted Transition System): A weighted *Transition System* wTS_i is a tuple $\left(\mathcal{Q}_i, q_i^0, \mathcal{A}_i, \rightarrow_i, w_i, \mathcal{AP}, L_i\right)$ where: (a) $\mathcal{Q}_i = \{q_i^{\ell_j}\}_{\forall \ell_j \in \mathcal{C}_m, \forall m \in \mathcal{S}_i}$ is the set of states, where a state $q_i^{\ell_j}$ indicates that robot i is at ℓ_j; (b) $q_i^0 \in \mathcal{Q}_i$ is the initial state of robot i; (c) $\mathcal{A}_i$ is a set of actions. The available actions at state $q_i^{\ell_j}$ are "wait" and "go to state $q_i^{\ell_k}$" for every k such that $(j, k) \in \mathcal{E}$;[1](d) $\rightarrow_i \subseteq \mathcal{Q}_i \times \mathcal{A}_i \times \mathcal{Q}_i$ is the transition relation; (e) $w_i : \mathcal{Q}_i \times \mathcal{Q}_i \to \mathbb{R}_+$ is a cost function that assigns weights/cost to each possible transition in wTS. These costs are associated with the distance between two states $q_i^{\ell_j}$ and $q_i^{\ell_k}$; (f) $\mathcal{AP}$ is the set of atomic propositions; and (h) $L_i : \mathcal{Q}_i \to 2^{\mathcal{AP}}$ is an observation/output relation giving the set of atomic propositions that are satisfied in a state.

In what follows we give definitions related to wTS_i, that we will use throughout the rest of the paper.

[1]Throughout the rest of the paper, for the sake of simplicity we assume that the mobility graph $\mathcal{G}$ is constructed so that there is a path from any location ℓ_j to any ℓ_k, where $q_i^{\ell_j} \in \mathcal{Q}_i$ and $q_i^{\ell_k} \in \mathcal{Q}_i$, that does not pass though a location $\ell_e \in \mathcal{C}_m$ for some $m \notin \mathcal{S}_i$, for all robots i. This assumption can be relaxed if we incorporate all communication points ℓ_j in the set of states $\mathcal{Q}_i$.

Definition 2.2 (Infinite Path): An *infinite path* τ_i of wTS_i is an infinite sequence of states, $\tau_i = \tau_i(1)\tau_i(2)\tau_{ij}(3)\dots$ such that $\tau_i(1) = q_i^0$, $\tau_i(n) \in \mathcal{Q}_i$, and $(\tau_i(n), a_i, \tau_i(n+1)) \in \rightarrow_i$, for some $a_i \in \mathcal{A}_i$, $\forall n$.[2]

Definition 2.3 (Composition): Composition of M infinite paths $\tau_m = \tau_m(1)\tau_m(2)\tau_m(3)\dots$, where $m \in \{1,\dots,M\}$, denoted by $\tau = \otimes_{\forall m}\tau_m$ is an infinite sequence of states defined as $\tau = \tau(1)\tau(2)\cdots = [\tau(n)]_{n=1}^{\infty}$, where $\tau(n) = (\tau_1(n), \tau_2(n), \dots, \tau_M(n))$.

Definition 2.4 (Cost): The cost of an infinite path τ_i of a wTS_i is $J_i(\tau_i) = \sum_{n=1}^{\infty} w_i(\tau_i(n), \tau_i(n+1))$. Similarly, the cost incurred by the composition of N infinite paths τ_i denoted by $\tau = \otimes_{\forall i\in\{1,2,\dots,N\}}\tau_i$ is given by

$$J(\tau) = \sum_{n=0}^{\infty}\sum_{i=1}^{N} w_i(\tau_i(n), \tau_i(n+1)), \qquad (2)$$

Definition 2.5 (Projection): For an infinite path $\tau = \tau(1)\tau(2)\tau(3)\dots$, we denote by $\Pi|_{\text{wTS}_i}\tau$ its projection onto wTS_i, which is obtained by erasing all states in τ that do not belong to $\mathcal{Q}_i$.

Definition 2.6 (Trace of infinite path): The trace of an infinite path $\tau_i = \tau_i(1)\tau_i(2)\tau_i(3)\dots$ of a wTS wTS_i, denoted by $\texttt{trace}(\tau_i)$, is an infinite word that is determined by the sequence of atomic propositions that are true in the states along τ_i, i.e., $\texttt{trace}(\tau_i) = L_i(\tau_i(1))L_i(\tau_i(2))\dots$.

Definition 2.7 (Motion Plan): Given an LTL formula ϕ, a wTS wTS_i both defined over the set of atomic propositions $\mathcal{AP}$, an infinite path τ_i of wTS_i is called *motion plan* if and only if $\texttt{trace}(\tau_i) \in \texttt{Words}(\phi)$, where $\texttt{Words}(\phi) = \{\sigma \in (2^{\mathcal{AP}})^{\omega} | \sigma \models \phi\}$ is defined as the set of words $\sigma \in (2^{\mathcal{AP}})^{\omega}$ that satisfy the LTL ϕ and $\models \subseteq (2^{\mathcal{AP}}) \times \phi$ is the satisfaction relation. The relation $\texttt{trace}(\tau_i) \in \texttt{Words}(\phi)$ is equivalently denoted by $\tau_i \models \phi$.

The problem we address in this paper can be stated as:

Problem 1: Given any initial configuration of the robots in the mobility graph $\mathcal{G}$ determine motion plans τ_i for all robots i that satisfy global LTL expression given in (1), i.e., communication graph $\mathcal{G}_c$ is connected over time and minimize the total distance traveled by robots captured by the objective function (2).

III. Intermittent Communication Control

To solve Problem 1, known centralized model checking techniques can be employed, that typically rely on a discretized abstraction of the environment captured by a wTS and the construction of a synchronized product system among all robots of the network. As a result, such approaches are resource demanding and scale poorly with the size of the network. Therefore, a distributed solution is preferred whereby discrete high-level motion plans for every robot can be computed locally across the network. For this purpose, notice first that although the global LTL formula (1) is not decomposable with respect to robots, it can be decomposed in local LTL formulas $\phi_{\mathcal{T}_m}$ associated with a robot team $\mathcal{T}_m$, which are coupled with each other by the conjunction operator $\wedge$. Specifically, we can write $\phi = \wedge_{m\in\mathcal{V}_{\mathcal{T}}}\phi_{\mathcal{T}_m}$, where $\phi_{\mathcal{T}_m}$ is defined as

$$\phi_{\mathcal{T}_m} = \square\Diamond\left(\vee_{\ell_j\in\mathcal{C}_m}(\wedge_{\forall i\in\mathcal{T}_m}\pi_i^{\ell_j})\right), \qquad (3)$$

and forces all robots $i \in \mathcal{T}_m$ to meet infinitely often at at least one rendezvous point $\ell_j \in \mathcal{C}_m$.

Given the decomposition of ϕ into local LTL formulas $\phi_{\mathcal{T}_m}$, every robot $i \in \mathcal{T}_m$ needs to develop motion plans τ_i so that the composition of plans τ_i, $\forall i \in \mathcal{T}_m$ denoted by $\tau_{\mathcal{T}_m} = \otimes_{i\in\mathcal{T}_m}\tau_i$ satisfies the local LTL expression $\phi_{\mathcal{T}_m}$, for all $m \in \mathcal{S}_i$. In this way, we can ensure that the composition of τ_i, $\forall i \in \{1, 2\dots, N\}$, satisfies the global LTL expression (1), since all local LTL expressions $\phi_{\mathcal{T}_m}$ are satisfied.

Motion plans $\tau_{\mathcal{T}_m} \models \phi_{\mathcal{T}_m}$, $\forall m \in \mathcal{V}_{\mathcal{T}}$, can be constructed using existing tools from model checking theory [5], [6]. However, notice that constructing plans $\tau_{\mathcal{T}_m}$ and $\tau_{\mathcal{T}_n}$, $\forall n \in \mathcal{N}_{\mathcal{T}_m}$ independently cannot ensure that the robots' behavior in the workspace will satisfy the global LTL formula (1). The reason is that the local LTL formulas $\phi_{\mathcal{T}_m}$ in (3) are not independent from the local LTL expressions $\phi_{\mathcal{T}_n}$ for which it holds $n \in \mathcal{N}_{\mathcal{T}_m}$, since they are coupled by robots' state in their respective transition systems. For instance, assume that $n \in \mathcal{N}_{\mathcal{T}_m}$ and that robot i belongs to teams $\mathcal{T}_m$ and $\mathcal{T}_n$. Then robot i is responsible for communicating with the other robots that belong to teams $\mathcal{T}_m$ and $\mathcal{T}_n$ at a communication point in $\mathcal{C}_m$ and $\mathcal{C}_n$, respectively. This equivalently implies that the LTL expressions $\phi_{\mathcal{T}_m}$ and $\phi_{\mathcal{T}_n}$ are coupled due to robot i through the atomic propositions $\pi_i^{\ell_j}$, $\forall \ell_j \in \mathcal{C}_m$ and $\pi_i^{\ell_k}$, $\forall \ell_k \in \mathcal{C}_n$. Consequently, generating plans $\tau_{\mathcal{T}_m} \models \phi_{\mathcal{T}_m}$ that ignore the LTL expressions $\phi_{\mathcal{T}_n}$, $\forall n \in \mathcal{N}_{\mathcal{T}_m}$ may result in conflicting robot behaviors, since the projection of motion plans $\tau_{\mathcal{T}_m}$ and $\tau_{\mathcal{T}_n}$ onto wTS_i may result in two different motion plans τ_i for a robot $i \in \mathcal{T}_m \cap \mathcal{T}_n$, $n \in \mathcal{N}_{\mathcal{T}_m}$. This means that cases where a robot i needs to be in more than one states in wTS_i simultaneously may occur.

To circumvent these issues, we propose a distributed algorithm in Section III-B that implements free-of-conflict discrete motion plans τ_i, for all robots i, so that the global LTL expression ϕ is satisfied. These motion plans will be constructed based on the prefix parts of motion plans $\tau_{\mathcal{T}_m} \models \phi_{\mathcal{T}_m}$, constructed in Section III-A, for all $m \in \mathcal{S}_i$.

A. Optimal Automata-based Model Checking

Given an LTL formula $\phi_{\mathcal{T}_m}$ and the wTS_i of all robots $i \in \mathcal{T}_m$ a motion plan $\tau_{\mathcal{T}_m} \models \phi_{\mathcal{T}_m}$ can be implemented using existing automata-based model checking methods [5], [6]. First the *weighted Product Transition System* (wPTS) $\text{wPTS}_{\mathcal{T}_m}$ is constructed, which essentially captures all the possible combinations of robots' states in their respective wTS_i, $\forall i \in \mathcal{T}_m$ and is defined as follows:

Definition 3.1 (Product Transition System): Given $|\mathcal{T}_m|$ weighted Transition Systems $\text{wTS}_{i_k} = (\mathcal{Q}_{i_k}, q_{i_k}^0, \mathcal{A}_{i_k}, \rightarrow_{i_k}, w_{i_k}, \mathcal{AP}, L_{i_k})$, where $i_k \in \mathcal{T}_m$, $k = 1, 2, \dots, |\mathcal{T}_m|$, the *weighted Product Transition System* $\text{wPTS}_{\mathcal{T}_m} = \text{wTS}_{i_1} \otimes \text{wTS}_{i_2} \otimes \cdots \otimes \text{wTS}_{i_{|\mathcal{T}_m|}}$ is

[2] A *finite path* of wTS_i can be defined accordingly. The only difference with the infinite path is that a finite path is defined by an finite sequence of states of wTS_i.

a tuple $\left(\mathcal{Q}_{\mathcal{T}_m}, q^0_{\mathcal{T}_m}, \mathcal{A}_{\mathcal{T}_m}, \longrightarrow_{\mathcal{T}_m}, w_{\mathcal{T}_m}, \mathcal{AP}, L_{\mathcal{T}_m}\right)$ where: (a)$\mathcal{Q}_{\mathcal{T}_m} = \mathcal{Q}_{i_1} \times \mathcal{Q}_{i_2} \times \cdots \times \mathcal{Q}_{i_{|\mathcal{T}_m|}}$ is the set of states; (b) $q^0_{\mathcal{T}_m} = (q^0_{i_1}, q^0_{i_2}, \ldots, q^0_{i_{|\mathcal{T}_m|}}) \in \mathcal{Q}_{\mathcal{T}_m}$ is the initial state; (c)$\mathcal{A}_{\mathcal{T}_m} = \mathcal{A}_{i_1} \times \mathcal{A}_{i_2} \times \cdots \times \mathcal{A}_{i_{|\mathcal{T}_m|}}$ is a set of actions; (d)$\longrightarrow_{\mathcal{T}_m} \subseteq \mathcal{Q}_{\mathcal{T}_m} \times \mathcal{A}_{\mathcal{T}_m} \times \mathcal{Q}_{\mathcal{T}_m}$ is the transition relation defined by the rule[3] $\dfrac{\wedge_{\forall i_k}\left(q_{i_k} \xrightarrow{a_{i_k}}_{i_k} q^{'}_{i_k}\right)}{q_{\mathcal{T}_m} \xrightarrow{a_{\mathcal{T}_m}=\left(a_{i_1},\ldots,a_{i_{|\mathcal{T}_m|}}\right)}_{\mathcal{T}_m} q^{'}_{\mathcal{T}_m}}$;[4] (e)$w_{\mathcal{T}_m}(q_{\mathcal{T}_m}, q^{'}_{\mathcal{T}_m}) = \sum_{k=1}^{\mathcal{T}_m} w_i(\Pi|_{\text{wTS}_{i_k}} q_{\mathcal{T}_m}, \Pi|_{\text{wTS}_{i_k}} q^{'}_{\mathcal{T}_m})$; (f)$\mathcal{AP}$ is the set of atomic propositions; and,(h)$L_{\mathcal{T}_m} = \bigcup_{i_k \in \mathcal{T}_m} L_{i_k}$ is an observation/output relation giving the set of atomic propositions that are satisfied at a state.

Next the LTL formula $\phi_{\mathcal{T}_m}$ is translated into a Nondeterministic Büchi Automaton (NBA) over $2^{\mathcal{AP}}$ denoted by $B_{\mathcal{T}_m}$ [17], which is defined as follows:

Definition 3.2: A *Nondeterministic Büchi Automaton* (NBA) $B_{\mathcal{T}_m}$ over $2^{\mathcal{AP}}$ is defined by the tuple $B_{\mathcal{T}_m} = \left(\mathcal{Q}_{B_{\mathcal{T}_m}}, \mathcal{Q}^0_{B_{\mathcal{T}_m}}, 2^{\mathcal{AP}}, \rightarrow_{B_{\mathcal{T}_m}}, \mathcal{F}_{B_{\mathcal{T}_m}}\right)$ where: (a) $\mathcal{Q}_{B_{\mathcal{T}_m}}$ is the set of states; (b) $\mathcal{Q}^0_{B_{\mathcal{T}_m}} \subseteq \mathcal{Q}_{B_{\mathcal{T}_m}}$ is a set of initial states; (c) $\Sigma = 2^{\mathcal{AP}}$ is an alphabet; (d) $\rightarrow_{B_{\mathcal{T}_m}} \subseteq \mathcal{Q}_{B_{\mathcal{T}_m}} \times \Sigma \times \mathcal{Q}_{B_{\mathcal{T}_m}}$ is the transition relation; and (e) $\mathcal{F}_{B_{\mathcal{T}_m}} \subseteq \mathcal{Q}_{B_{\mathcal{T}_m}}$ is a set of accepting/final states.

Once the wPTS $\text{wPTS}_{\mathcal{T}_m}$ and the NBA $B_{\mathcal{T}_m}$ that corresponds to the LTL $\phi_{\mathcal{T}_m}$ are constructed, a motion plan $\tau_{\mathcal{T}_m} \models \phi_{\mathcal{T}_m}$ can be found by checking the non-emptiness of the language of the *Product Büchi Automaton* (PBA) $P_{\mathcal{T}_m} = \text{wPTS}_{\mathcal{T}_m} \otimes B_{\mathcal{T}_m}$ [5], which is defined as follows:

Definition 3.3 (Product Büchi Automaton): Given the product transition system $\text{wPTS}_{\mathcal{T}_m} = \left(\mathcal{Q}_{\mathcal{T}_m}, q^0_{\mathcal{T}_m}, \mathcal{A}_{\mathcal{T}_m}, \longrightarrow_{\mathcal{T}_m}, w_{\mathcal{T}_m}, \mathcal{AP}, L_{\mathcal{T}_m}\right)$ and the NBA $B_{\mathcal{T}_m} = \left(\mathcal{Q}_{B_{\mathcal{T}_m}}, \mathcal{Q}^0_{B_{\mathcal{T}_m}}, 2^{\mathcal{AP}}, \rightarrow_{B_{\mathcal{T}_m}}, \mathcal{F}_{B_{\mathcal{T}_m}}\right)$, the *Product Büchi Automaton* (PBA) $P_{\mathcal{T}_m} = \text{wPTS}_{\mathcal{T}_m} \otimes B_{\mathcal{T}_m}$ is a tuple $\left(\mathcal{Q}_{P_{\mathcal{T}_m}}, \mathcal{Q}^0_{P_{\mathcal{T}_m}}, \longrightarrow_{P_{\mathcal{T}_m}}, \mathcal{F}_{P_{\mathcal{T}_m}}\right)$ where: (a) $\mathcal{Q}_{P_{\mathcal{T}_m}} = \mathcal{Q}_{\mathcal{T}_m} \times \mathcal{Q}_{B_{\mathcal{T}_m}}$ is the set of states; (b) $\mathcal{Q}^0_{P_{\mathcal{T}_m}} = q^0_{\mathcal{T}_m} \times \mathcal{Q}^0_{B_{\mathcal{T}_m}}$ is a set of initial states; (c) $\longrightarrow_{P_{\mathcal{T}_m}} \subseteq \mathcal{Q}_{P_{\mathcal{T}_m}} \times \mathcal{A}_{\mathcal{T}_m} \times 2^{\mathcal{AP}} \times \mathcal{Q}_{P_{\mathcal{T}_m}}$ is the transition relation defined by the rule: $\dfrac{\left(q_{\mathcal{T}_m} \xrightarrow{a_{\mathcal{T}_m}} q^{'}_{\mathcal{T}_m}\right) \wedge \left(q_{B_{\mathcal{T}_m}} \xrightarrow{L_{\mathcal{T}_m}\left(q^{'}_{\mathcal{T}_m}\right)} q^{'}_{B_{\mathcal{T}_m}}\right)}{q_{P_{\mathcal{T}_m}} = \left(q_{\mathcal{T}_m}, q_{B_{\mathcal{T}_m}}\right) \xrightarrow{a_{\mathcal{T}_m}}_{P_{\mathcal{T}_m}} q^{'}_{P_{\mathcal{T}_m}} = \left(q^{'}_{\mathcal{T}_m}, q^{'}_{B_{\mathcal{T}_m}}\right)}$; (d) $\mathcal{F}_{P_{\mathcal{T}_m}} = \mathcal{Q}_{\mathcal{T}_m} \times \mathcal{F}_{B_{\mathcal{T}_m}}$ is a set of accepting/final states.

To check the non-emptiness of the language of $P_{\mathcal{T}_m}$ denoted by $\mathcal{L}_{P_{\mathcal{T}_m}} = \texttt{trace}(\text{wPTS}_{\mathcal{T}_m}) \cap \mathcal{L}_{B_{\mathcal{T}_m}}$ and to find the motion plan that both satisfies $\phi_{\mathcal{T}_m}$ and at the same time minimizes $J(\tau_{\mathcal{T}_m})$, we can employ existing model checking methods that are based on graph search algorithms; see, e.g., [18], [19]. Such motion plans can be written in a prefix-suffix structure $\tau_{\mathcal{T}_m} = \tau^{\text{pre}}_{\mathcal{T}_m}[\tau^{\text{suf}}_{\mathcal{T}_m}]^{\omega}$, where the prefix part $\tau^{\text{pre}}_{\mathcal{T}_m}$ is executed only once and the suffix part $\tau^{\text{suf}}_{\mathcal{T}_m}$ is repeated infinitely. In principle, in these approaches the PBA is viewed as a weighted directed graph with weights assigned on each edge that are inherited by the function $w_{\mathcal{T}_m}$. Then finding the shortest path from an initial state to a final state and projecting this path onto $\text{wPTS}_{\mathcal{T}_m}$ results in the prefix part $\tau^{\text{pre}}_{\mathcal{T}_m}$. The suffix part $\tau^{\text{suf}}_{\mathcal{T}_m}$ is constructed similarly by computing the shortest cycle around that final state.

[3]The notation of this rule is along the lines of the notation used in [5]. In particular, it means that if the proposition above the solid line is true, then so does the proposition below the solid line.

[4]The state $q_{\mathcal{T}_m}$ stands for the state $\left(q_{i_1}, \ldots, q_{i_{|\mathcal{T}_m|}}\right) \in \mathcal{Q}_{\mathcal{T}_m}$. The state $q^{'}_{\mathcal{T}_m}$ is defined accordingly.

B. Conflict Resolution Coordination

As discussed in the beginning of Section III, constructing motion plans $\tau_{\mathcal{T}_m}$ for all $m \in \mathcal{S}_i$, for all robots i can result in conflicting robot behaviors. To overcome this issue, we propose a distributed algorithm that resolves any conflicts in the robot behavior introduced by the motion plans $\tau_{\mathcal{T}_m}$ and constructs free-of-conflicts motion plans τ_i for all robots i using the prefix parts $\tau^{\text{pre}}_{\mathcal{T}_m}$ constructed in Section III-A. The general form of these motion plans is

$$\begin{aligned}\tau_i =& \tau_i(1)\tau_i(2)\cdots = [\tau_i(n)]_{n=1}^{\infty}\\ =& \Big[X \ldots X\Pi|_{\text{wTS}_i} p^k_{\mathcal{T}_{m_1}} X \ldots X\Pi|_{\text{wTS}_i} p^k_{\mathcal{T}_{m_j}}\\ & X \ldots X\Pi|_{\text{wTS}_i} p^k_{\mathcal{T}_{m_{|\mathcal{S}_i|}}} X \ldots X\Big]_{k=1}^{\infty} = \left[p^k_i\right]_{k=1}^{\infty}, \quad (4)\end{aligned}$$

such that $\tau_i(1) = q^0_i$. In (4), $p^k_{\mathcal{T}_{m_j}}$ is a finite path of $\text{wPTS}_{\mathcal{T}_m}$, where $m_j \in \mathcal{S}_i$ and $j \in \{1, \ldots, |\mathcal{S}_i|\}$. Also, X stands for a finite path in which robot i waits at its current state in wTS_i. In (4), the concatenation of the paths X and $p^k_{\mathcal{T}_{m_j}}$, $\forall m_j \in \mathcal{S}_i$ gives rise to the finite path p^k_i. Hereafter, the e-th finite path in p^k_i is denoted by $p^{k,e}_i$ where $1 \leq e \leq \ell$, where ℓ stands for the number of finite paths that appear in p^k_i. The parameter ℓ is *a priori* selected to be $\ell = \max\{d_{\mathcal{T}_m}\}_{m=1}^{M} + 1$ for all robots, where $d_{\mathcal{T}_m}$ denotes the degree of vertex m in the graph $\mathcal{G}_{\mathcal{T}}$. This particular choice for the parameter ℓ ensures the construction of free-of-conflict motion plans, as it will shown in Proposition 3.6. Moreover, we denote by $p^k_i(n)$ the n-th state in the finite path p^k_i, e.g., $p^1_i(1) = p^{1,1}_i(1) = \tau_i(1) = q^0_i$, by construction of τ_i. The same notation extends to the infinite path τ_i.

In what follows we first describe the construction of the finite paths $p^k_{\mathcal{T}_{m_j}}$ and then we show how these finite paths are ordered in p^k_i giving rise to a free-of-conflict motion plan τ_i. First, for the finite path $p^k_{\mathcal{T}_{m_j}}$ it holds that $p^k_{\mathcal{T}_{m_j}} = \tau^{\text{pre},k}_{\mathcal{T}_{m_j}}$, where $\tau^{\text{pre},k}_{\mathcal{T}_{m_j}}$ is the prefix part of the motion plan $\tau^k_{\mathcal{T}_{m_j}} \models \phi_{\mathcal{T}_{m_j}}$ constructed as per Section III-A. The index k is introduced in $\tau^{\text{pre},k}_{\mathcal{T}_{m_j}}$ to point out that the prefix part is recomputed as the index k, introduced in (4) changes. The reason it needs to be recomputed is because the initial state of $\text{wPTS}_{\mathcal{T}_{m_j}}$ changes as k changes. Particularly, the initial state of $\text{wPTS}_{\mathcal{T}_{m_j}}$ is the state $q^{0,k}_{\mathcal{T}_{m_j}} = p^k_{\mathcal{T}_{m_j}}(1)$. The state $\Pi|_{\text{wTS}_i} p^k_{\mathcal{T}_{m_j}}(1)$ is selected so that all sub-paths $\Pi|_{\text{wTS}_i} p^k_{\mathcal{T}_{m_j}}$ in p^k_i chain up consistently. Hence, we select the state $\Pi|_{\text{wTS}_i} p^k_{\mathcal{T}_{m_j}}(1)$ to be the final state of the previous sub-path that appeared in p^k_i. Consequently, the state $p^k_i(1)$ coincides with the the final state of the finite path p^{k-1}_i. If

$k = 1$, then $p_i^k(1)$ refers to the initial position of robot i in the workspace.

The finite paths p_i^k are constructed sequentially across the nodes $\ell_j \in \mathcal{V}$, as follows. Let $\mathcal{S} = \{\ell_1, \dots, \ell_j, \dots\}$ be an ordered sequence of the nodes in the mobility graph $\mathcal{G}$, so that consecutive nodes (communication points) ℓ_j, ℓ_e in $\mathcal{S}$ are associated with teams $\mathcal{T}_n$ and $\mathcal{T}_m$, respectively, that belong to neighboring nodes in the graph $\mathcal{G}_\mathcal{T}$, i.e., $\ell_e \in \mathcal{C}_m$, $\ell_j \in \mathcal{C}_n$ and $m \neq n$ and $m \in \mathcal{N}_{\mathcal{T}_n}$. We assume that $\mathcal{S}$ is known by all robots and that every robot i is initially located at the first communication point $\ell_e \in \mathcal{C}_m$, $m \in \mathcal{S}_i$ that appears in $\mathcal{S}$. Assume that paths have been constructed for all nodes in $\mathcal{S}$ that precede $\ell_e \in \mathcal{C}_m$ and that currently all robots $i \in \mathcal{T}_m$ are located at node ℓ_e and coordinate to construct the paths p_i^k. Since the mobility graph $\mathcal{G}$ is connected consecutive nodes in $\mathcal{S}$ are connected by a path in $\mathcal{G}$, this means that there is at least one robot $j \in \mathcal{T}_n \cap \mathcal{T}_m$, $n \in \mathcal{N}_{\mathcal{T}_m}$, which previously constructed its path p_j^k by placing at its $n_j^{\mathcal{T}_m}$-th entry of the finite path $\Pi|_{\text{wTS}_j} p^k_{\mathcal{T}_m}$, i.e., $p_j^{k,n_j^{\mathcal{T}_m}} = \Pi|_{\text{wTS}_j} p^k_{\mathcal{T}_m}$. Then robot i constructs the path p_i^k based on three rules. According to the first rule, the path $\Pi|_{\text{wTS}_i} p^k_{\mathcal{T}_m}$ will be placed at the $n_i^{\mathcal{T}_m}$-th entry, which is selected to be equal to $n_j^{\mathcal{T}_m}$, which is common for all robots $j \in \mathcal{T}_m$ [line 1, Alg. 1]. This ensures that robot i and all other robots $j \in \mathcal{T}_m$ will meet at a communication point that belongs to $\mathcal{C}_m$, as it will be shown in Proposition 3.8. The next step is to place the paths $\Pi|_{\text{wTS}_i} p^k_{\mathcal{T}_g}$ for all $g \in \mathcal{S}_i \backslash \{m\}$, at the $n_i^{\mathcal{T}_g}$-th entry of p_i^k. The index $n_i^{\mathcal{T}_g}$ will be determined by one of the two following rules. Specifically, according to the second rule if there exist robots $j \in \mathcal{N}_i \cap \mathcal{T}_g$ that have already constructed the paths p_j^k, then the index $n_i^{\mathcal{T}_g}$ is selected to be equal to $n_j^{\mathcal{T}_g}$, which is common for all $j \in \mathcal{T}_g$ [line 4, Alg. 1]. Otherwise, according to the third rule the path $\Pi|_{\text{wTS}_i} p^k_{\mathcal{T}_g}$ can be placed at any free entry of p_i^k indexed by $n_i^{\mathcal{T}_g}$, provided that the $n_i^{\mathcal{T}_g}$-th entry of all paths p_b^k of robots $b \in \mathcal{N}_i$ that have already been constructed does not contain states $\Pi|_{\text{wTS}_b} p_{\mathcal{T}_h}(k)$ with $h \in \mathcal{N}_{\mathcal{T}_g}$ [line 6, Alg. 1]. To highlight the role of this rule assume that $h \in \mathcal{N}_{\mathcal{T}_g}$. Then this means that there exists at least one robot $r \in \mathcal{T}_h \cap \mathcal{T}_g$. Then notice that without the third rule [line 5], at a subsequent iteration of this procedure, robot $r \in \mathcal{T}_h \cap \mathcal{T}_g$ would have to place the paths $\Pi|_{\text{wTS}_r} p^k_{\mathcal{T}_g}$ and $\Pi|_{\text{wTS}_r} p^k_{\mathcal{T}_h}$ at a common entry of p_r^k, i.e., $n_r^{\mathcal{T}_g} = n_r^{\mathcal{T}_h}$, due to the two previous rules and, therefore, a conflicting behavior for robot r would occur. In all the remaining entries of p_i^k, Xs are placed [line 8, Alg. 1].[5] This procedure is repeated until all robots $i \in \mathcal{T}_m$ have constructed their respective paths p_i^k. Once this happens, all robots $i \in \mathcal{T}_m$ depart from node $\ell_e \in \mathcal{C}_m$ and travel to the next communication point $\ell_c \in \mathcal{C}_v$, that appears in $\mathcal{S}$ that satisfies $v \in \mathcal{S}_i$ [line 9, Alg. 1]. At that point, all robots associated with the next communication point in $\mathcal{S}$ are present at that node, and can coordinate to compute their respective paths, as before. The procedure is repeated sequentially over the nodes in $\mathcal{S}$ until all robots have computed their paths.

When all robots have constructed their finite paths, they exchange a set of indices denoted by $\mathcal{X}_i$ that collects the indices n_i^X at which $p_i^{k,n_i^X} = X$. If there exist paths $p_i^{k,n_i^X} = X$, for some $n_i^X \in \bigcap_{\forall i} \mathcal{X}_i$, they are discarded, since in these paths all robots i wait at their current states. Also, notice that in general, the finite paths $p_i^{k,e}$ for some $e \in \{1, \dots, \ell\}$ may have different lengths across the robots i. Consequently, this implies that two robots i, j that belong to a team $\mathcal{T}_m$ may start executing the finite paths $\Pi|_{\text{wTS}_i} p^k_{\mathcal{T}_m}$ and $\Pi|_{\text{wTS}_j} p^k_{\mathcal{T}_m}$ at different time instants, assuming that the robots pick synchronously their next states in their transition systems. Avoiding such a case is crucial to ensure intermittent communication within team $\mathcal{T}_m$, as it will be shown in Proposition 3.8. Therefore, given any index e we can introduce states at the end of the finite paths $p_i^{k,e}$ where the robots wait in their current states so that all finite paths $p_i^{k,e}$, for all robots i, have the same length. Communication between the robots in the last two stages of the algorithm can happen in the order defined by $\mathcal{S}$, as before.

Remark 3.4: Note that communication according to $\mathcal{S}$ is very predictable and inefficient as it, e.g., does not allow for simultaneous meetings at the nodes of $\mathcal{G}$. For these reasons it is only used to construct conflict-free motion plans that allow for much more efficient intermittent communication between robots.

Remark 3.5 (Optimality): To construct the motion plans τ_i in (4), we first decouple the global LTL formula (1) into local LTL expressions $\phi_{\mathcal{T}_m}$. Then, we construct finite paths $p^k_{\mathcal{T}_{m_j}} = \tau^{\text{pre},k}_{\mathcal{T}_m}$, where $\tau^k_{\mathcal{T}_m} \models \phi_{\mathcal{T}_m}$, for all $m_j \in \mathcal{S}_i$, and their concatenation (Algorithm 1) gives rise to the motion plans τ_i. In principle, the paths $p^k_{\mathcal{T}_{m_j}}$ connect the current configuration of robots $i \in \mathcal{T}_{m_j}$ to the closest meeting point $\ell_e \in \mathcal{C}_{m_j}$ where all robots in $\mathcal{T}_{m_j}$ meet. Therefore, while these paths minimize the total distance traveled between the current and next meeting points, they do not optimize the infinite horizon cost function (2). Consequently, the proposed solution is suboptimal. To obtain an optimal solution that minimizes (2), standard model checking techniques can be applied to the global product system, that are known to be computationally expensive. Suboptimality here is a consequence of problem decomposition.

C. Correctness of the Proposed Algorithm

In this section, we show that the composition of the discrete motion plans τ_i generated by Algorithm 1 satisfies the global LTL expression (1), i.e., that the network is connected over time. To prove this result, we need first to show that Algorithm 1 can develop non-conflicting motion plans τ_i, for which we have the following two results. Proposition 3.6 can be proved by following the steps of the proof of

[5] If $\ell_j = \ell_1 \in \mathcal{C}_q$, for some $q \in \mathcal{V}_\mathcal{T}$, then initially, a randomly selected robot $j \in \mathcal{T}_q$ creates arbitrarily its path p_j^k by placing the paths $\Pi|_{\text{wTS}_j} p^k_{\mathcal{T}_q}$ at the $n_j^{\mathcal{T}_q}$-th entry of p_j^k, for all $q \in \mathcal{S}_j$. Then the procedure previously described follows. Moreover, depending on the structure of the graph $\mathcal{G}_\mathcal{T}$ it is possible that a communication point $\ell_j \in \mathcal{C}_m$ appears more than once in $\mathcal{S}$. In this case, robots $i \in \mathcal{T}_m$ construct the finite paths p_i^k only the first time that ℓ_j appears in $\mathcal{S}$.

Algorithm 1 Construction of motion plans $\tau_i = [p_i^k]_{k=1}^{\infty}$ at node $\ell_e \in \mathcal{C}_m$

Require: Already constructed finite paths p_j^k of robots $j \in \mathcal{N}_i$;

Require: All robots in $\mathcal{T}_m$ are located at node $\ell_e \in \mathcal{C}_m$;

$p_i^{k,n_i^{\mathcal{T}_m}} := \Pi|_{\text{wTS}_i} p_{\mathcal{T}_m}^k$, $n_i^{\mathcal{T}_m} = n_j^{\mathcal{T}_m},\ \forall j \in \mathcal{T}_m$;

for $g \in \mathcal{S}_i \setminus \{m\}$ **do**

  if there exist constructed paths p_j^k, $j \in \mathcal{N}_i \cap \mathcal{T}_g$ **then**

    $p_i^{k,n_i^{\mathcal{T}_g}} := \Pi|_{\text{wTS}_i} p_{\mathcal{T}_g}^k$, $n_i^{\mathcal{T}_g} = n_j^{\mathcal{T}_g},\ \forall j \in \mathcal{T}_g$;

  else

    $p_i^{k,n_i^{\mathcal{T}_g}} := \Pi|_{\text{wTS}_i} p_{\mathcal{T}_g}^k$ provided either $p_j^{k,n_i^{\mathcal{T}_g}} = X$, or $p_j^{k,n_i^{\mathcal{T}_g}} = \Pi|_{\text{wTS}_j} p_{\mathcal{T}_h}^k$ with $h \notin \mathcal{N}_{\mathcal{T}_g},\ \forall j \in \mathcal{N}_i$;

  end if

  Put Xs in the remaining entries;

  Transmit path p_i^k to a robot in $\mathcal{T}_m$ that has not constructed its motion plan. If there are not such robots, all robots $i \in \mathcal{T}_m$ depart from node $\ell_e \in \mathcal{C}_m$;

end for

Proposition 3.2 in [14] for the graph $\mathcal{G}_{\mathcal{T}}$. Proposition 3.7 holds by construction of the finite paths p_i^k and its proof, which is omitted due to space limitations, is along the lines of the proof of Proposition 3.3 in [14].

Proposition 3.6: Algorithm 1 can always construct finite paths p_i^k that consist of ℓ finite paths $p_i^{k,e}$ where $\ell \leq \max\{d_{\mathcal{T}_m}\}_{m=1}^{M} + 1$.

Proposition 3.6 shows also that the finite paths p_i^k and consequently, the motion plans τ_i depend on the node degree of graph $\mathcal{G}_{\mathcal{T}}$, and not on the size of the network.

Proposition 3.7: Algorithm 1 generates *admissible* discrete motion plans τ_i, i.e., motion plans that are free of conflicts and satisfy the transition rule $\rightarrow_i$ defined in Definition 2.1.

Proposition 3.8: The composition of motion plans τ_i generated by Algorithm 1 satisfies the global LTL expression (1), i.e., connectivity of the robot network is ensured over time, infinitely often.

Proof: To prove this result, it suffices to show that $\tau_{\mathcal{T}_m,*} \models \phi_{\mathcal{T}_m}$, for all teams $m \in \mathcal{V}_{\mathcal{T}}$, where $\phi_{\mathcal{T}_m}$ is defined in (3) and $\tau_{\mathcal{T}_m,*} = \otimes_{\forall i \in \mathcal{T}_m} \tau_i$, where the motion plans τ_i are generated by Algorithm 1. Equivalently, according to Definition 2.7 it suffices to show that

$$\texttt{trace}(\tau_{\mathcal{T}_m,*}) \in \texttt{Words}(\phi_{\mathcal{T}_m}). \tag{5}$$

First, assume that the prefix part $\tau_{\mathcal{T}_m}^{\text{pre},k}$ of the motion plan $\tau_{\mathcal{T}_m}^k = \tau_{\mathcal{T}_m}^{\text{pre},k}[\tau_{\mathcal{T}_m}^{\text{suf},k}]^{\omega} \models \phi_{\mathcal{T}_m}$ has been constructed, for all k, using a standard model checking method as described in Section III-A.[6] By construction of the motion plans τ_i, it holds that there is an index $n_i^{\mathcal{T}_m}$, $1 \leq n_i^{\mathcal{T}_m} \leq \ell$, such that $p_i^{k,n_i^{\mathcal{T}_m}} = \Pi|_{\text{wTS}_i} p_{\mathcal{T}_m}^k$ and an index $e_i^{\mathcal{T}_m}$, which are common for all robots $i \in \mathcal{T}_m$, such that $p_i^k(e_i^{\mathcal{T}_m}) = \Pi|_{\text{wTS}_i} p_{\mathcal{T}_m}^k(1)$ for all $i \in \mathcal{T}_m$. Consequently, in the finite path $\pi_{\mathcal{T}_m,*}^k = \otimes_{\forall i \in \mathcal{T}_m} p_i^k$ there are indices $n^{\mathcal{T}_m}$ and $e^{\mathcal{T}_m}$ such that $\pi_{\mathcal{T}_m,*}^{k,n^{\mathcal{T}_m}} = p_{\mathcal{T}_m}^k$ and $\pi_{\mathcal{T}_m,*}^k(e^{\mathcal{T}_m}) = p_{\mathcal{T}_m}^k(1)$ for all k. Hence, we conclude that all robots $i \in \mathcal{T}_m$ will start executing the finite path $p_{\mathcal{T}_m}^k$, simultaneously.

Using the above observation, we examine the properties of the finite word that is generated by the team $\mathcal{T}_m$ when the finite path $p_{\mathcal{T}_m}^k$ is executed. Let $w_{\mathcal{T}_m}^k \in (2^{\mathcal{AP}})^*$ be a finite word defined as $w_{\mathcal{T}_m}^k = \texttt{trace}(\tau_{\mathcal{T}_m}^{\text{pre},k})$ or, equivalently, by construction of the finite paths $p_{\mathcal{T}_m}^k$, $w_{\mathcal{T}_m}^k = \texttt{trace}(p_{\mathcal{T}_m}^k)$. Since $\tau_{\mathcal{T}_m}^k \models \phi_{\mathcal{T}_m}$ and by construction of the prefix part $\tau_{\mathcal{T}_m}^{\text{pre},k}$ for all k, we have that the infinite word $\sigma_{\mathcal{T}_m} = w_{\mathcal{T}_m}^0 w_{\mathcal{T}_m}^1 w_{\mathcal{T}_m}^2 \cdots \in (2^{\mathcal{AP}})^{\omega}$ satisfies

$$\sigma_{\mathcal{T}_m} \in \texttt{Words}(\phi_{\mathcal{T}_m}). \tag{6}$$

Next, notice that between the execution of the finite paths $p_{\mathcal{T}_m}^k$ and $p_{\mathcal{T}_m}^{k+1}$ robots $i \in \mathcal{T}_m$ will traverse through some states of their respective wTSs that are determined by $\tau_{\mathcal{T}_m,*}$ until the state $p_{\mathcal{T}_m}^{k+1}(1)$ is reached. In other words, the team of robots $\mathcal{T}_m$ does not execute consecutively the finite paths $p_{\mathcal{T}_m}^k$ and $p_{\mathcal{T}_m}^{k+1}$. Now, we need to show that these intermediate transitions that robots make are admissible in $\text{wPTS}_{\mathcal{T}_m}$ and do not violate $\phi_{\mathcal{T}_m}$. These transitions are admissible as implied by Proposition 3.7. Also, they cannot violate $\phi_{\mathcal{T}_m}$, since the LTL expressions $\phi_{\mathcal{T}_m}$ in (3), for all $m \in \mathcal{V}_{\mathcal{T}}$, do not include the negation operator $\neg$ in front of the atomic propositions $\pi_i^{\ell_j}$ defined in Section II-A. Therefore, this equivalently means that as the robots $i \in \mathcal{T}_m$ execute the motion plan $\tau_{\mathcal{T}_m,*}$ they will eventually pass through all states determined by the finite paths $p_{\mathcal{T}_m}^k$, for all k, without violating $\phi_{\mathcal{T}_m}$. Consequently, this means that as the robots $i \in \mathcal{T}_m$ move according to the motion plan $\tau_{\mathcal{T}_m,*}$, the generated trace $\texttt{trace}(\tau_{\mathcal{T}_m,*})$ will certainly include the atomic propositions that are included in $\sigma_{\mathcal{T}_m}$; see also Definition 2.6. In other words, the sequence of atomic propositions $\sigma_{\mathcal{T}_m}$ is a subsequence of $\texttt{trace}(\tau_{\mathcal{T}_m,*})$ while the additional atomic propositions that exist in $\texttt{trace}(\tau_{\mathcal{T}_m,*})$ cannot violate $\phi_{\mathcal{T}_m}$, as previously discussed. Therefore, due to (6), we conclude that (5) holds for all $m \in \mathcal{V}_{\mathcal{T}}$, which completes the proof. ■

In general, the motion plans $\tau_i = [\tau_i(n)]_{n=1}^{\infty}$ defined in (4) are infinite paths of wTS_i, since the finite paths p_i^k need to be updated for every $k \in \mathbb{N}$. Therefore, in practice they are hard to implement and manipulate. In the following proposition, we show that the motion plans τ_i constructed by Algorithm 1 have a finite representation and they can be expressed in a prefix-suffix structure, where the prefix part τ_i^{pre} is traversed only once and the suffix part τ_i^{suf} is repeated infinitely.

Proposition 3.9: Algorithm 1 generates discrete motion plans τ_i for all robots i in a prefix-suffix structure, i.e., $\tau_i = \tau_i^{\text{pre}} \left[\tau_i^{\text{suf}}\right]^{\omega} = [p_i^1 \dots p_i^{k_p - 1}][p_i^{k_p} \dots p_i^{k_s}]^{\omega}$.

Proof: To show this result it suffices to show that there is an index k_s, such that for all $k \geq k_s$, the finite paths p_i^k are repeated and, therefore, they do not need to be recomputed. First, notice that such an index k_s is common for all motion

[6] Note that, throughout this proof, the only difference between the plans $\tau_{\mathcal{T}_m,*}$ and $\tau_{\mathcal{T}_m}$ is that the first one has been derived through composing all non-conflicting motion plans τ_i, $\forall i \in \mathcal{T}_m$ generated by Algorithm 1, while the second one is the plan for all robots $i \in \mathcal{T}_m$ computed as described in Section III-A.

plans τ_i. To illustrate this point assume that there is a robot $i \in \mathcal{T}_m$ such that $\Pi|_{\text{wTS}_i} p^k_{\mathcal{T}_m} \neq \Pi|_{\text{wTS}_i} p^{k+1}_{\mathcal{T}_m}$ and, therefore, $p^k_i \neq p^{k+1}_i$. Since the paths $p^k_{\mathcal{T}_m}$ are computed collectively by all robots $i \in \mathcal{T}_m$, the fact that $\Pi|_{\text{wTS}_i} p^k_{\mathcal{T}_m} \neq \Pi|_{\text{wTS}_i} p^{k+1}_{\mathcal{T}_m}$ means that there may be another robot $j \in \mathcal{T}_m$ such that $\Pi|_{\text{wTS}_j} p^k_{\mathcal{T}_m} \neq \Pi|_{\text{wTS}_j} p^{k+1}_{\mathcal{T}_m}$ and, consequently, $p^k_j \neq p^{k+1}_j$, which may be propagated to all other robots in the network, since graph $\mathcal{G}_\mathcal{T}$ is connected. Therefore, the index k_s is common for all robots.

In what follows, we show the existence of an index k_s defined above. By definition the finite paths p^k_i are a concatenation of ℓ finite paths, where the e-th finite path in p^k_i is denoted by $p^{k,e}_i$. Also, by construction we have that $p^{k,e}_i \neq X$ is constructed by computing a shortest path in a graph associated with a PBA $P_{\mathcal{T}_m}$ from an initial to a final state in $P_{\mathcal{T}_m}$. Notice that the set of final states $\mathcal{F}_{P_{\mathcal{T}_m}}$ in $P_{\mathcal{T}_m}$ remain the same for all k, which is not the case for the set of initial states $\mathcal{Q}^0_{P_{\mathcal{T}_m}} = q^{0,k}_{\mathcal{T}_m} \times \mathcal{Q}^0_{B_{\mathcal{T}_m}}$. The reason is that the initial state $q^{0,k}_{\mathcal{T}_m}$ of wPTS$_{\mathcal{T}_m}$ may change over k and, specifically, we have that $q^{0,k}_{\mathcal{T}_m} = p^k_{\mathcal{T}_m}(1) = (p^{k,e}_{i_1}(1), \ldots, p^{k,e}_{i_{|\mathcal{T}_m|}}(1))$, where the construction of state $p^k_{\mathcal{T}_m}(1)$ was presented in Section III-A. Since the number of possible combinations of robots' states in their respective wTS is finite, we have that there is a finite number of possible initial states $q^{0,k}_{\mathcal{T}_m}$ for all $m \in \mathcal{V}_\mathcal{T}$ and, consequently, a finite number of possible initial states $q^{0,k} = (p^{k,e}_1(1), \ldots, p^{k,e}_N(1))$. Therefore, for any $e \in \{1, \ldots, \ell\}$ there are always two indices k^e_1 and k^e_2, $k^e_2 > k^e_1$ such that $p^{k^e_1,e}_i(1) = p^{k^e_2,e}_i(1)$, for all robots i, since otherwise that would mean that there are infinite number of possible $q^{0,k} = (p^{k,e}_1(1), \ldots, p^{k,e}_N(1))$ for any $e \in \{1, \ldots, \ell\}$. Since $p^{k^e_1,e}_i(1) = p^{k^e_2,e}_i(1)$ and the distances between communication points remains the same for all k, we have that $p^{k^e_1,e}_i = p^{k^e_2,e}_i$. Then, the finite paths $p^{k,e}_i$, for all $e \in \{1, 2, \ldots, \ell\}$, for all robots i, and for all $k \geq k^e_2$ have already been constructed at k that satisfies $k^e_1 \leq k \leq k^e_2$; for instance, for $k = k^e_2 + 1$ it holds that $p^{k^e_2+1,e}_i = p^{k^e_1+1,e}_i$. Let $e^* = \text{argmin}_e \{k^e_2\}^\ell_{e=1} = \text{argmin}_e \{k^e_1\}^\ell_{e=1}$. Then, similarly, we have that the finite paths $p^{k,e}_i$, for all $e \in \{1, 2, \ldots, \ell\}$, for all robots i, and for all $k \geq k^{e*}_2$ have already been constructed at k that satisfies $k^{e^*}_1 \leq k \leq k^{e^*}_2$. Consequently, this means that the finite paths p^k_i for all $k \geq k^{e^*}_2$ have already been constructed as well at k that satisfies $k^{e^*}_1 \leq k \leq k^{e^*}_2$; for instance $p^{k^{e^*}_1}_i = p^{k^{e^*}_2}_i$ and $p^{k^{e^*}_1+1}_i = p^{k^{e^*}_2+1}_i$. Hence, the motion plans τ_i can be written in a prefix-suffix structure $\tau_i = \tau^{\text{pre}}_i \left[\tau^{\text{suf}}_i\right]^\omega$ where $\tau^{\text{pre}}_i = [p^1_i \ldots p^{k^{e^*}_1 - 1}_i]$ and $\tau^{\text{suf}}_i = [p^{k^{e^*}_1}_i \ldots p^{k^{e^*}_2 - 1}_i p^{k^{e^*}_2}_i]$, i.e., $k_p = k^{e^*}_1$ and $k_s = k^{e^*}_2$, which completes the proof. ■

IV. Asynchronous Intermittent Communication

In section III, we showed that if all robots i pick synchronously their next states in wTS$_i$ according to the motion plans τ_i, then the LTL expression (1) is satisfied. In this section, we show that the generated motion plans can be executed asynchronously, as well, by appropriately introducing delays in the continuous-time execution of τ_i. We omit a formal proof of this result due to space limitations, and instead validate the proposed asynchronous scheme through numerical simulations in Section V.

Due to the asynchronous execution of the controllers, the motion plans τ_i can be written as in (4) replacing the indices n and k with n_i and k_i, respectively, which allows us to model the situation where the robots pick asynchronously their next states in wTS$_i$. In the asynchronous execution of the infinite paths τ_i robot i moves from state $\tau_i(n_i - 1) \in \mathcal{Q}_i$ to $\tau_i(n_i) \in \mathcal{Q}_i$ according to a continuous-time motion controller $\mathbf{u}_i(t) \in \mathbb{R}^n$ that belongs to the tangent space of γ_{ij} at $\mathbf{x}_i(t)$. Without loss of generality, assume that $i \in \mathcal{T}_m$ and $\tau_i(n_i) = q^{\ell_e}_i$, for some $\ell_e \in \mathcal{C}_m$. When robot i arrives at state $\tau_i(n_i)$ it checks if $\tau_j(n_i) = \Pi|_{\text{wTS}_j} p_{\mathcal{T}_m}(f_i) = q^{\ell_e}_j$ for all $j \in \mathcal{T}_m$. If this is the case, then robot i waits at node ℓ_e until all other robots $j \in \mathcal{T}_m$ are present there. When this happens, or if there is at least one robot $j \in \mathcal{T}_m$ such that $\tau_j(n_i) \neq \Pi|_{\text{wTS}_j} p_{\mathcal{T}_m}(f_i) = q^{\ell_e}_j$, then robot i moves towards the next state $\tau_i(n_i + 1)$.

V. Simulation Studies

In this section, a simulation study is provided that illustrates our approach for a network of $N = 5$ robots that move along the edges of the mobility graph with $L = 20$ communication points as shown in Figure 2. The network is divided in $M = 5$ teams which are $\mathcal{T}_1 = \{1, 2\}$, $\mathcal{T}_2 = \{2, 3\}$, $\mathcal{T}_3 = \{3, 4\}$, $\mathcal{T}_4 = \{2, 4, 5\}$, and $\mathcal{T}_5 = \{1, 5\}$ and, therefore, the graph $\mathcal{G}_\mathcal{T}$ is as shown in Figure 1(b). Also, the mobility graph is constructed so that there is a path γ_{ij} from any node ℓ_i to any other node ℓ_j. Therefore, the finite paths $p^k_{\mathcal{T}_m}$ constructed as per Section III-A have the form $p^k_{\mathcal{T}_m} = q^{0,k}_{\mathcal{T}_m} q_{\mathcal{T}_m}$, where $q^{0,k}_{\mathcal{T}_m}$ is the initial state of wPTS$_{\mathcal{T}_m}$ constructed as defined in Section III-B and $q_{\mathcal{T}_m} = (q^{\ell_j}_{i_1} \ldots q^{\ell_j}_{i_{|\mathcal{T}_m|}})$ is a state where all robots of team $\mathcal{T}_m$ are located at a common meeting point $\ell_j \in \mathcal{C}_m$. The motion plans τ_i generated by Algorithm 1 have the following structure:

$$
\begin{aligned}
\tau_1 &= \left[p^{k_1}_1\right]^\infty_{k_1=1} = \left[\Pi|_{\text{wTS}_1} \tau^{k_1}_{\mathcal{T}_1} \Pi|_{\text{wTS}_1} \tau^{k_1}_{\mathcal{T}_5} X\right]^\infty_{k_1=1}, \\
\tau_2 &= \left[p^{k_2}_2\right]^\infty_{k_2=1} = \left[\Pi|_{\text{wTS}_2} \tau^{k_2}_{\mathcal{T}_1} \Pi|_{\text{wTS}_2} \tau^{k_2}_{\mathcal{T}_2} \Pi|_{\text{wTS}_2} \tau^{k_2}_{\mathcal{T}_4}\right]^\infty_{k_2=1}, \\
\tau_3 &= \left[p^{k_3}_3\right]^\infty_{k_3=1} = \left[\Pi|_{\text{wTS}_3} \tau^{k_3}_{\mathcal{T}_3} \Pi|_{\text{wTS}_3} \tau^{k_3}_{\mathcal{T}_2} X\right]^\infty_{k_3=1}, \\
\tau_4 &= \left[p^{k_4}_4\right]^\infty_{k_4=1} = \left[\Pi|_{\text{wTS}_4} \tau^{k_4}_{\mathcal{T}_3} X \Pi|_{\text{wTS}_4} \tau^{k_4}_{\mathcal{T}_4}\right]^\infty_{k_4=1}, \\
\tau_5 &= \left[p^{k_5}_5\right]^\infty_{k_5=1} = \left[X \Pi|_{\text{wTS}_5} \tau^{k_5}_{\mathcal{T}_5} \Pi|_{\text{wTS}_5} \tau^{k_5}_{\mathcal{T}_4}\right]^\infty_{k_5=1},
\end{aligned}
$$

which can be written in a prefix-suffix structure with $k_s = 2$ and $k_p = 1$, where the indices k_s and k_p are defined in Proposition 3.9. The motion plans τ_i defined above are depicted in Figure 2. Notice that the distances between any two meeting points vary across $\mathcal{G}$ and so do the robots' velocities and, therefore, robots pick asynchronously their next states in wTS$_i$. Consequently, this results in waiting times for every robot $i \in \mathcal{T}_m$ at the meeting points $\ell_j \in \mathcal{C}_m$, which are non-integer multiples of each other, for all $m \in$

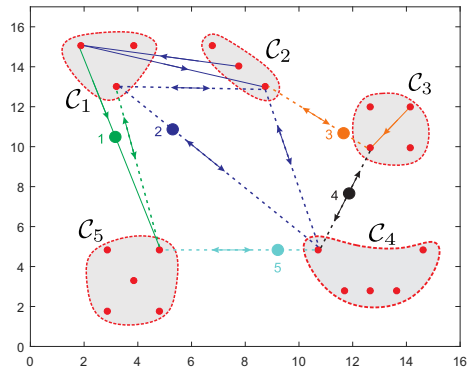

Fig. 2. Intermittent communication of $N = 5$ robots moving along the edges of an underlying mobility graph. Red dots represent communication points in each set $\mathcal{C}_m$ while edges between any two communication points exist (not shown). Straight lines depict the robots' trajectories as determined by motion plans τ_i. The suffix part of τ_i is represented by dashed lines for every robot i while solid lines depict a part of prefix structure, i.e., a path that connects robots' initial states to the respective suffix structure. The prefix part of robots 4 and 5 coincides with their respective suffix part and, therefore, there are no corresponding solid lines for them.

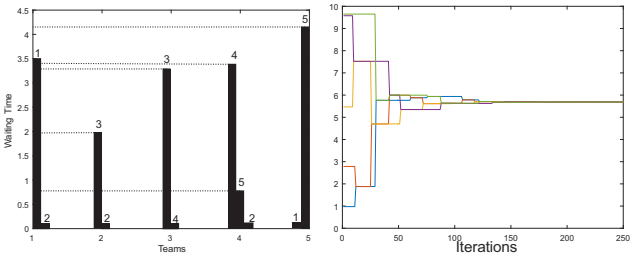

(a) Waiting Time (b) Consensus

Fig. 3. Figure 3(a) depicts the waiting time of every robot i at a meeting point $\ell_j \in \mathcal{C}_m$, for all $m \in \mathcal{S}_i$ during a single execution of τ_i^{suf}. Figure 3(b) illustrates the consensus of numbers $v_i(t)$.

$\mathcal{V}_\mathcal{T}$. This illustrated in Figure 3(a), where, e.g., in team $\mathcal{T}_4$, robots 2, 4, and 5 wait at a meeting point $\ell_j \in \mathcal{C}_4$ for 0, 3.4, and 0.8 time units, respectively. Observe also in Figure 3(a) that robot 2 never waits at any meeting point.

To illustrate that under the proposed motion plans connectivity is ensured over time, we implement a consensus algorithm over the dynamic network $\mathcal{G}_c$. Specifically, we assume that initially robots generate a random number $v_i(t_0)$ and when all robots $i \in \mathcal{T}_m$ meet at $\ell_j \in \mathcal{C}_m$ they perform the following consensus update: $v_i(t) = \frac{1}{|\mathcal{T}_m|} \sum_{e \in \mathcal{T}_m} v_e(t)$. Figure 3(b) shows that eventually all robots reach a consensus on the numbers $v_i(t)$. Note also that applying existing LTL-based planning would result in PBA constructed for all robots with $\Pi_{i=1}^N |\mathcal{Q}_i||\mathcal{Q}_B| = 45360|\mathcal{Q}_B|$ states, where $\mathcal{Q}_B$ corresponds to the state-space of the NBA associated with the global LTL expression in (1), which is hard to manipulate in practice. This issue becomes more severe as the size of the network increases.

VI. Conclusion

In this paper we considered the problem of controlling intermittent communication in mobile robot networks. We assumed that robots move along the edges of a mobility graph and they can communicate only when they meet at its nodes, which gave rise to a dynamic communication network. The network was defined to be connected over time if communication takes place at the rendezvous points infinitely often which was captured by an LTL formula. Then this LTL expression was approximately decomposed into local LTL expressions which were assigned to robots. To avoid conflicting robot behaviors that could occur due to this approximate decomposition, we developed a distributed conflict resolution scheme that generated free-of-conflicts discrete motion plans for every robot that ensured connectivity over time, infinitely often and minimized the distance traveled by the robots. We also showed that the generated motion plans can be executed asynchronously by introducing delays in their continuous-time execution.

References

[1] M. M. Zavlanos and G. J. Pappas, "Potential fields for maintaining connectivity of mobile networks," *IEEE Transactions on Robotics*,, vol. 23, no. 4, pp. 812–816, 2007.

[2] L. Sabattini, N. Chopra, and C. Secchi, "Decentralized connectivity maintenance for cooperative control of mobile robotic systems," *The International Journal of Robotics Research*, vol. 32, no. 12, pp. 1411–1423, 2013.

[3] M. Zavlanos and G. Pappas, "Distributed connectivity control of mobile networks," *IEEE Transactions on Robotics*, vol. 24, no. 6, pp. 1416–1428, 2008.

[4] M. Zavlanos, M. Egerstedt, and G. Pappas, "Graph theoretic connectivity control of mobile robot networks," *Proc. of the IEEE*, vol. 99, no. 9, pp. 1525–1540, 2011.

[5] C. Baier and J.-P. Katoen, *Principles of model checking*. MIT press Cambridge, 2008, vol. 26202649.

[6] E. M. Clarke, O. Grumberg, and D. Peled, *Model checking*. MIT press, 1999.

[7] S. G. Loizou and K. J. Kyriakopoulos, "Automatic synthesis of multi-agent motion tasks based on ltl specifications," in *43rd IEEE Conference on Decision and Control (CDC)*, vol. 1, The Bahamas, December 2004, pp. 153–158.

[8] H. Kress-Gazit, G. E. Fainekos, and G. J. Pappas, "Temporal-logic-based reactive mission and motion planning," *IEEE Transactions on Robotics*, vol. 25, no. 6, pp. 1370–1381, 2009.

[9] M. Guo, J. Tumova, and D. V. Dimarogonas, "Cooperative decentralized multi-agent control under local ltl tasks and connectivity constraints," in *53rd Conference on Decision and Control (CDC)*, Los Angeles, CA, USA, December 2014, pp. 75–80.

[10] Y. Chen, X. C. Ding, and C. Belta, "Synthesis of distributed control and communication schemes from global ltl specifications," in *50th IEEE Conference on Decision and Control and European Control Conference*, Orlando, FL, USA, December 2011, pp. 2718–2723.

[11] M. Kloetzer and C. Belta, "Distributed implementations of global temporal logic motion specifications," in *IEEE International Conference on Robotics and Automation*, Pasadena, CA, USA, May 2008, pp. 393–398.

[12] A. Ulusoy, S. L. Smith, X. C. Ding, C. Belta, and D. Rus, "Optimality and robustness in multi-robot path planning with temporal logic constraints," *The International Journal of Robotics Research*, vol. 32, no. 8, pp. 889–911, 2013.

[13] Y. Kantaros and M. M. Zavlanos, "Intermittent connectivity control in mobile robot networks," in *49th Asilomar Conference on Signals, Systems and Computers*, Pacific Grove, CA, USA, November, 2015, pp. 1125–1129.

[14] ——, "A distributed ltl-based approach for intermittent communication in mobile robot networks," in *American Control Conference*, Boston, MA, USA, July, 2016, pp. 5557–5562.

[15] M. M. Zavlanos, "Synchronous rendezvous of very-low-range wireless agents," in *49th IEEE Conference on Decision and Control (CDC)*, Atlanta, GA, USA, December 2010, pp. 4740–4745.

[16] G. Hollinger and S. Singh, "Multi-robot coordination with periodic connectivity," in *IEEE International Conference on Robotics and Automation (ICRA)*, Anchorage, Alaska, May 2010, pp. 4457–4462.

[17] M. Y. Vardi and P. Wolper, “An automata-theoretic approach to automatic program verification,” in *1st Symposium in Logic in Computer Science (LICS)*. IEEE Computer Society, 1986.

[18] M. Guo, K. H. Johansson, and D. V. Dimarogonas, “Motion and action planning under ltl specifications using navigation functions and action description language,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Tokyo Big Sight, Japan, November 2013, pp. 240–245.

[19] M. Guo and D. V. Dimarogonas, “Reconfiguration in motion planning of single-and multi-agent systems under infeasible local ltl specifications,” in *IEEE 52nd Annual Conference on Decision and Control (CDC)*, Florence, December 2013, pp. 2758–2763.